NoFAQ: Synthesizing Command Repairs from Examples

Loris D'Antoni
University of Wisconsin-Madison
loris@cs.wisc.edu

Rishabh Singh
Microsoft Research
risin@microsoft.com

Michael Vaughn
University of Wisconsin-Madison
vaughn@cs.wisc.edu

arXiv:1608.08219v1 [cs.PL] 29 Aug 2016

Abstract

Command-line tools are confusing and hard to use for novice programmers due to their cryptic error messages and lack of documentation. Novice users often resort to online help-forums for finding corrections to their buggy commands, but have a hard time in searching precisely for posts that are relevant to their problem and then applying the suggested solutions to their buggy command.

We present a tool NoFAQ that uses a set of rules to suggest possible fixes when users write buggy commands that trigger commonly occurring errors. The rules are expressed in a language called Fixit and each rule pattern-matches against the user's buggy command and the corresponding error message, and uses these inputs to produce a possible fixed command. Our main contribution is an algorithm based on lazy VSA for synthesizing Fixit rules from examples of buggy and repaired commands. The algorithm allows users to add new rules in NoFAQ without having to manually encode them. We present the evaluation of NoFAQ on 92 benchmark problems and show that NoFAQ is able to instantly synthesize rules for 81 benchmark problems in real time using just 2 to 5 input-output examples for each rule.

1. Introduction

Command-Line Interfaces (CLI) let users interact with a computing system by writing sequences of commands. CLIs are especially popular amongst advanced computer users, who use them to perform small routine tasks such as committing a file to a repository with version control, installing software packages, compiling source code, finding and searching for files etc. Even though this mode of interaction is getting replaced by more natural graphical user interfaces, CLIs are still routinely used for most scripting tasks in Unix and Mac OS. Even the Windows operating system now officially provides complex command-line interfaces with products such as Windows Powershell.

Since command-line interactions often require complex parameters and flag settings for specifying the desired intent, non-expert users find CLIs challenging to use [3]. Moreover, after entering an incorrect input command, the user has to deal with cryptic errors that are hard to decipher by just looking at the verbose text-based documentation of the commands. For these reasons, users typically resort to online help-forums for finding corrections to their buggy commands. Unfortunately, this can also be problematic as users need to precisely search for posts that relate to the issues with their commands and then transform the suggested solutions to apply them in their context. Sometimes users also need to create a new post if no relevant post exists (or can be found), and then need to wait for hours or days to obtain a solution to their problem.

What about common errors? Recently, a tool TheFxxx[1] was developed for automatically addressing common errors when working with a CLI. If after typing a command a user receives an error message, TheFxxx uses a set of hard-coded rules to suggest possible fixes to the user's command. Each rule pattern-matches against the input command and the error message, and uses these inputs to produce a possible fixed command. Typical fixes include adding missing flags, creating a missing directory, or changing file extensions. TheFxxx has become extremely popular and, on GitHub, it has already been starred by more than 20,000 users and has been forked more than 1000 times. Despite its success, the tool also has a main limitation: to add a new rule a developer first needs to understand the syntax and precise semantics of TheFxxx and then manually hard-code the rule into the tool. Due to this complexity, newly added rules have at times caused non-terminating or unexpected behaviors[2].

Synthesizing repairs from examples Inspired by the success and limitations of TheFxxx, we built NoFAQ (No more Frequently Asked Questions), a tool for automatically addressing common errors in CLIs. NoFAQ also uses a set of rules for fixing common errors, but it differs from TheFxxx in the following key aspects:

1. Rules are encoded in a declarative domain-specific language (DSL) called Fixit.
2. To add new rules, users only provide examples of buggy and repaired commands and NoFAQ automatically synthesizes the desired Fixit rules that are consistent with the given examples.

We envision NoFAQ system being used by non-developers and end-users, who can easily extend the system by providing new examples of fixes. The long-term goal of the system is to learn from examples obtained from shell histories of thousands of users in an unsupervised manner. Although a developer can write similar rules in a system like TheFxxx manually, there are two main challenges with doing so. First, it is not feasible to easily add thousands of rules as end-users generally do not have contributor access to TheFxxx's source code. Second, even for developers, writing correct rules is difficult and error-prone because of the complexity of the string manipulations needed to perform the command fixes. In fact, TheFxxx only consists of less than 100 rules in a little over 1 year, since adding new ones is a fairly complex task.

The Fixit DSL for encoding fix rules is inspired by the types of rules appearing in TheFxxx and by common command repairs requested by users on help-forums. A Fixit rule first uses pattern matching and unification to match the command and error message, and then applies a fix transformation if the match succeeds. The transformations consist of substring and append functions on strings present in the command and error message.

We present an algorithm that efficiently synthesizes Fixit rules that are consistent with a given set of input-output examples using a Version-space Algebra (VSA). VSA-based synthesis techniques are used to succinctly represent the set of all expressions that are

[1] We decided to censor the name of the tool. The tool can be found at http://bit.ly/CmdCorrection.

[2] http://bit.ly/1j7zxOr and http://bit.ly/1YgngXJ.

consistent with a set of examples [8, 11]. Even though existing VSA data structures can represent an exponential number of FIXIT rules in polynomial space, this space can still be quite large. To address this problem we introduce *lazy version-space algebra*. Given a set of examples, our algorithm maintains a lazy representation of only a subset of the FIXIT rules that are consistent with the examples. The rules missing from the version-space are only enumerated when necessary—i.e., when a new input-output example can only be accounted by adding a FIXIT rule that is not already present in the version-space. Because of the careful design of FIXIT, our synthesis algorithm has a polynomial time complexity. In contrast, existing VSA-based synthesis techniques for string transformations require exponential time [8]. The polynomial time complexity is crucial for our synthesis algorithm to scale to a large number of fix examples. Since different examples may refer to different target rules, we propose a strategy to partition the input examples into groups of examples corresponding to individual rules. We then use the lazy VSA algorithm to learn the FIXIT rules for each partition.

We evaluate the synthesis algorithm implemented in NOFAQ on 92 benchmark problems obtained from both THEFUXXX (76) and online help-forums (16). NOFAQ is able to learn the repair rules for 81 of the buggy commands in these benchmark problems from only 2 to 5 input-output examples each.

Contributions summary:

1. FIXIT, a declarative domain-specific language for encoding rules that map a command and an error message to possible fixed commands (§ 3).
2. A sound and complete polynomial time synthesis algorithm based on lazy version-space algebra for synthesizing FIXIT rules from input-output examples (§ 4).
3. An analysis of the formal properties of the FIXIT language and its synthesis algorithm (§ 5).
4. A qualitative and quantitative evaluation of the synthesis algorithm on 92 benchmarks obtained from both THEFUXXX and online help-forums (§ 6.2).

2. Motivating examples

We first present the main ideas behind NOFAQ using some concrete examples. All the example rules presented in this section are actual rules appearing in THEFUXXX system.

2.1 Adding missing file extension

Java programmers, in particular novice ones, are likely to encounter this error when they accidentally pass a class name instead of a source code file to the `javac` compiler.

cmd1:	javac Employee
err1:	Class names, 'Employee', are only accepted if annotation processing is explicitly requested

Often, this error is provided by the `javac` compiler when it is invoked on a file that does not have the proper `.java` extension. A seasoned programmer would immediately recognize the problem and add the extension `.java` at the end of the input file.

fix1:	javac Employee.java

On the other hand, a novice programmer will search the web in the hope of finding a way to address the error and understand how to apply it to their setting. The goal of NOFAQ is to automatically synthesize simple fixes like this one from examples provided by experienced users and use the synthesized fixes to help novice users who encounter similar errors. For example, let's say that a skilled developer provides another triple of the following form.

cmd2:	javac Pair
err2:	Class names, 'Pair', are only accepted if annotation processing is explicitly requested
fix2:	javac Pair.java

Using these two examples NOFAQ will synthesize the following fix rule.

match [STR(`javac`), VAR-MATCH$(1, \varepsilon, \varepsilon)$]
and [STR(`Class`), STR(`names,`), VAR-MATCH$(2,$ `‘,’`$,)$
STR(`are`), STR(`only`), STR(`accepted`), STR(`if`),
STR(`annotation`), STR(`processing`), STR(`is`),
STR(`explicitly`), STR(`requested`)]
$\rightarrow$ [FSTR(`javac`), SUB-LR$(0, 0, \varepsilon,$ `.java`, VAR(1))]

The first part of the rule (i.e., up to the symbol $\rightarrow$) pattern-matches against the command and the error message and binds the input strings to the corresponding variables. The variables are then used by the second part of the rule to produce the output. In this case the SUB-LR$(0, 0, \varepsilon,$ `.java`, VAR(1)) expression extracts the complete string associated with VAR(1) (a start index of 0 and an end index of 0 denotes the identity string extraction), and then prepends the string ε at the beginning of it, and appends the string `.java` at the end of it.

2.2 Extracting substrings

The following scenario is another common one for novice Java programmers.

cmd1:	java Run.java
err1:	Could not find or load main class Run.java
fix1:	java Run

Given this example and another similar one, NOFAQ synthesizes the following rule.

match [STR(`java`), VAR-MATCH$(1, \varepsilon,$ `.java`$)$]
and [STR(`Could`), STR(`not`), STR(`find`), STR(`or`),
STR(`load`), STR(`main`), STR(`class`),
VAR-MATCH$(2, \varepsilon,$ `.java`$)$]
$\rightarrow$ [FSTR(`javac`), SUB-LR$(0, -5, \varepsilon, \varepsilon,$ VAR(1))]

This rule extracts the substring of the input file name starting at index 0 and ending at the index -5 (5^{th} index from the end of the string) in order to remove the `.java` extension.

2.3 Extracting complex substrings

A user was trying to move a picture from one location to another but got the following error message.

cmd1:	mv photo.jpg Mary/summer12.jpg
err1:	can't rename 'photo.jpg': No such file or directory
fix1:	mkdir Mary `&&` mv photo.jpg Mary/summer12.jpg

The error is cryptic for novice command-line users and does not guide them towards the actual problem of the missing directory. Given this example and another similar one, NOFAQ synthesizes the following rule.

match [STR(mv), VAR-MATCH(1, ε, ε),
VAR-MATCH(2, ε, ε)]
and [STR(can't), STR(rename), VAR-MATCH(3, ', '), STR(No),
STR(such), STR(file), STR(or), STR(directory)]
→ [FSTR(mkdir), SUB-LR(0, CPOS(/, 1, 0), ε, &&, VAR(2))
FSTR(mv), SUB-LR(0, 0, ε, ε, VAR(1)),
SUB-LR(0, 0, ε, ε, VAR(2))]

The second expression in the output extracts the directory name—i.e., the substring that starts at index 0 and ends at the index of first occurrence of the character /. The rule also adds a string && at the end of the extracted string to pipe the two output commands.

3. The command repair language FIXIT

We now describe FIXIT, a domain-specific language for expressing repair rules. The syntax and semantics of FIXIT is presented in Figure 1 and Figure 2 respectively. The language FIXIT is designed to be expressive enough to capture most of the rules we found in THEFXXX and in online help-forums, but at the same time concise enough to enable efficient learning from examples.

General structure Each FIXIT program is a rule of the form

$$\textbf{match } cmd \textbf{ and } err \rightarrow fix$$

that takes as input a command $\bar{s}_{cmd}$ and an error $\bar{s}_{err}$ and either produces a fixed command or the undefined value $\perp$. The inputs $\bar{s}_{cmd}$ and $\bar{s}_{err}$ are lists of strings that are obtained by extracting all the space-separated strings appearing in the input command and error message respectively. The output fix produced by the rule is also a list of strings. From now on, we assume that the inputs and outputs are lists of strings that do not contain space characters.

A rule has three components.

1. A list of match expressions $cmd = [m_1, \cdots, m_l]$ used to pattern match against the input command $\bar{s}_{cmd}$.
2. A list of match expressions $err = [m_1, \cdots, m_k]$ used to pattern match against the input error message $\bar{s}_{err}$.
3. A list of fix expressions $fix = [f_1, \cdots, f_m]$ used to produce the new fixed command.

Match expressions A match expression m is either of the form STR(s) denoting a constant string s or of the form VAR-MATCH(i, l, r). A VAR-MATCH(i, l, r) expression denotes a variable index i and requires the matched string to start with the prefix l and end with the suffix r. We assume that no two variable expressions appearing in the match expression have the same index. When a list of match expressions $[m_1, \cdots, m_l]$ is applied to a list of strings $\bar{s} = [s_1, \ldots, s_l]$ with the same length l, it generates a partial function $\sigma : \mathbb{N} \mapsto \Sigma^*$ that assigns variables appearing in the match expressions to the corresponding strings in the input. For example, evaluating

$$[\text{STR(mv)}, \text{VAR-MATCH}(1, \epsilon, \texttt{.jpg}), \text{VAR-MATCH}(2, \epsilon, \texttt{.jpg})]$$

on the list of strings [mv, a.jpg, b.jpg] produces the function σ such that $\sigma(1) =$ a.jpg and $\sigma(2) =$ a.jpg. On the other hand, evaluating it on [mv, a.png, b.jpg] yields $\perp$, as a.png does not match the required suffix in VAR-MATCH(1, ϵ, .jpg).

Fix expressions A fix expression f is either of the form FSTR(s) denoting the constant output string s, or of the form SUB-LR(p_L, p_R, l, r, VAR(i)) denoting a function that is applied to the string s_i matched by the variable VAR(i). This function outputs the string $l \cdot m \cdot r$, where $\cdot$ denotes the string concatenation operator and $m = \texttt{substr}(s, j_L, j_R)$ where j_L and j_R are the indices resulting from respectively evaluating the position expressions p_L and p_R on the string s. Here, given a string $s = a_0 \ldots a_n$, the expression $\texttt{substr}(s, j_L, j_R)$ denotes the string $a_{j_L} \ldots a_{j_R - 1}$ if $j_L, j_R \leq n + 1$ and the undefined value $\perp$ otherwise. Notice that, unlike previous VSA-based languages [8], FIXIT does not allow binary recursive concatenation; this is one of the key features that enables polynomial time synthesis.

Fix rule	r	:=	**match** cmd **and** $err \rightarrow fix$
Input cmd	cmd	:=	$[m_1, \cdots, m_l]$
Input error	err	:=	$[m_1, \cdots, m_m]$
Output cmd	fix	:=	$[f_1, \cdots, f_n]$
Match expr	m	:=	STR(s)
		|	VAR-MATCH(i, l, r)
Fix expr	f	:=	FSTR(s)
		|	SUB-LR(p_L, p_R, l, r, VAR(i))
Pos expr	p	:=	IPOS(k)
		|	CPOS(c, k, δ)

Variables and constants:
s, s_l, s_r : *string* $\quad$ i, k, δ : *integer*
c : *character*

Figure 1: Syntax of the rule description language FIXIT.

Positions expressions A position expression p can be one of the following types of expressions.

- A constant position expression IPOS(k), which denotes the index k if k is positive and the index $|s| - k$ if k is negative. If $k = 0$, this expression evaluates to 0 when evaluated for p_L (i.e., the starting index of the substring) and to $|s|$ when evaluated for p_R, where $|s|$ denotes the length of the string s. For example, in the function SUB-LR(IPOS(0), IPOS(0), ε, ε, VAR(1)), where $\sigma(1) =$ File the first constant position evaluates to the index 0, while the second constant position evaluates to the index |File| = 4.
- A symbolic position expression CPOS(c, k, δ), which denotes the result of applying an offset δ to the index of the k-th occurrence of the character c in s if k is positive, and the result of applying an offset δ to the index of the k-th to last occurrence of the character c in s if k is negative. For example, given the string www.google.com, the expression CPOS(., 1, −2) denotes the index 2 (two positions before the first dot), while the expression CPOS(., −1, 2) denotes the index 12 (two positions after the last dot). This operator is novel and can express operations that are not supported by previous VSA-based work. In particular, FlashFill [8] only allows the extraction of the exact position of a character and not positions in its proximity. Despite this additional capability, FIXIT programs can be synthesized in polynomial time.

Comparison with FlashFill DSL At the top-level, FIXIT consists of match expressions over original command and error message, which perform pattern-matching and unification of variables with strings. This form of matching and unification is not expressible in FlashFill, so we cannot use it to learn the fix rules directly. However, we can use FlashFill as a subroutine to learn string transformations corresponding to expressions similar to SUB-LR expressions in NOFAQ. However, the FlashFill DSL has two major limitations: 1) No support for offsets from regular expression matches in computing position expressions (in contrast to FIXIT's CPOS(c, k, δ) operator), and 2) A finite hard-coded token set for regular expressions (e.g. no support for constant character tokens). Moreover, as described in subsection 4.3 and section 8, our SUB-LR operator yields a synthesis algorithm that operates in polynomial time in the

$$[\![\textbf{match } cmd \textbf{ and } err \ \rightarrow \ fix]\!](\bar{s}_{cmd}, \bar{s}_{err}) \ = \ \begin{cases} [\![fix]\!]_\sigma & \sigma = \texttt{unify}(cmd, \bar{s}_{cmd}) \cup \texttt{unify}(err, \bar{s}_{err}) \wedge \sigma \neq \bot \\ \bot & \text{otherwise} \end{cases}$$

$$\texttt{unify}([m_1, \cdots, m_l], [s_1, \cdots, s_l]) \ = \ \bigcup_{1 \leq j \leq l} \texttt{match-expr}(m_j, s_j),$$

$$\texttt{match-expr}(\textsc{Str}(s_1), s_2) \ = \ \begin{cases} \sigma_0 & \text{if } s_1 = s_2 \text{ and } \sigma_0 \text{ is the always undefined function} \\ \bot & \text{otherwise} \end{cases}$$

$$\texttt{match-expr}(\textsc{Var-Match}(i, l, r), s) \ = \ \begin{cases} \sigma & \text{if } \exists\, \delta \text{ s.t. } s = l\delta r \text{ and } \sigma(i) = s \text{ and } \sigma \text{ is undefined on every } j \neq i \\ \bot & \text{otherwise} \end{cases}$$

$$[\![[f_1, \cdots, f_n]]\!]_\sigma \ = \ [[\![f_1]\!]_\sigma, \cdots, [\![f_n]\!]_\sigma]$$

$$[\![\textsc{Fstr}(s)]\!]^{\textbf{fun}}_\sigma \ = \ s$$

$$[\![\textsc{Sub-lr}(p_L, p_R, l, r, \textsc{Var}(i))]\!]^{\textbf{fun}}_\sigma \ = \ l \cdot m \cdot r \qquad \text{if } (i, s) \in \sigma \text{ and } m = \texttt{substr}(s, [\![p_L]\!]^{\textbf{pos}}_{s,L}, [\![p_R]\!]^{\textbf{pos}}_{s,R})$$

$$[\![\textsc{Ipos}(k)]\!]^{\textbf{pos}}_{s,d} = \begin{cases} k & k > 0 \\ |s| + k & k < 0 \\ 0 & k = 0 \wedge d = L \\ |s| & k = 0 \wedge d = R \end{cases} \qquad [\![\textsc{Cpos}(c, k, \delta)]\!]^{\textbf{pos}}_{s,d} = \begin{cases} I[k-1] + \delta & k > 0 \text{ and } I = \texttt{indices}(s, c) \\ & \text{and } |I| \geq k \\ I[l-k] + \delta & k < 0 \text{ and } I = \texttt{indices}(s, c) \\ & \text{and } |I| + k \geq 0 \wedge l = |I| \\ \bot & \text{otherwise} \end{cases}$$

$$\texttt{indices}(s, c) = [i_0, \ldots, i_j] \qquad \text{where } \forall\, 0 \leq l < j : \ i_l < i_{l+1}, s[i_l] = c \text{ and } \forall j : s[j] = c \rightarrow \exists\, 0 \leq l < j : i_l = j$$

Figure 2: The semantics of the command repair DSL FIXIT.

number of examples, which enables the algorithm to scale to a large number of examples. Because of the recursive binary concatenate expressions in FlashFill, the DAG intersection based synthesis algorithm is exponential in the number of examples.

4. Synthesizing rules from examples

In this section we first describe our algorithm for synthesizing a single FIXIT rule from a set of examples of concrete command fixes. We then describe a multi-stage partitioning algorithm for learning multiple FIXIT programs from a large undifferentiated set of command repair examples.

The algorithm for learning a single FIXIT rule is described in Figure 8; it takes as input a list of examples $E = [e_1, \ldots, e_n]$ where each example e_i is a triple of the form $(\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})$ and outputs a symbolically represented set of FIXIT rules R consistent with E—i.e., for every rule $r \in R$ and example $e_i = (\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})$, the rule r outputs $\bar{s}_{fix}$ on the input $(\bar{s}_{cmd}, \bar{s}_{err})$.

The algorithm processes one input example at a time, and after processing the first i examples $E_i = [e_1, \ldots, e_i]$ the algorithm has computed a set of rules R_i consistent with E_i. At the end, the algorithm outputs one of the rules in R_n. We use $\bot$ to denote the undefined result. If at any point our algorithm returns $\bot$ it means that there is no FIXIT rule that is consistent with the given set of examples.

4.1 Symbolic representation of multiple rules

Since there can be exponentially many rules consistent with the input examples, we adopt a symbolic representation of the set R that is guaranteed to always have polynomial size. Our synthesis algorithm takes as input a list of examples E and produces as output a symbolic rule of the form **match** cmd **and** $err \ \rightarrow_s \ fixes$, where cmd and err are tuples of expressions that can consist of either constants or variables, and $fixes = [f_1, \ldots, f_m]$ is a list of expressions that symbolically represents a set of outputs that is consistent with the examples E. Formally, each f_i in $fixes$ is either a constant expression $\textsc{Fstr}(s)$ for some s, or a set of substring expressions $\{su_1, \ldots, su_k\}$, where each su_i is of the form $\textsc{Sub-lr}(p_L, p_R, l, r, \textsc{Var}(j))$. Intuitively, if we replace each set with one of the fix expressions it contains, we obtain a FIXIT rule. If each f_i contains k elements, this symbolic representation models k^n programs using an expression of size kn.

4.2 Lazy rule representation

The core element of our algorithm is a lazy representation of the rules that represents match and fix expressions as constants for as long as possible—i.e., until a new example shows that some parts of the rule cannot be constants. This helps reduce the number of variable expressions, which in turn reduces the number of substring expressions to be considered. We first illustrate the idea with a concrete example. Let's say that we are given the two examples shown in Figure 3a and 4a. After processing the first example, our algorithm synthesizes the FIXIT rule in Figure 5a in which every match expression and every fix expression is a constant. However, since we have only seen one example, we do not yet know whether some expression appearing in the match should actually be a variable match expression or whether some element in the fix expression should actually be a function of some variable. The main idea is that any of these possibilities can still be "recovered" when a new example shows that indeed a variable is needed. Using this idea, we maintain each expression as a constant until a new example shows that some expression cannot actually be a constant.

This is exactly what happens when processing the input example in Figure 4a. At this point in order to find a rule that is consistent with both examples we need to introduce a variable match as the second expression of the command match, and some function application as the second element of the fix. To do so, our algorithm applies the following operations to the previously computed rule.

1. All match expressions that cannot be constants are "promoted" to variable match expressions (making sure that all variable names are unique), which match on the longest shared prefix and suffix of all previously seen values at that position. The following table illustrates the idea for the case in which we try

cmd1: java Run.java
err1: Could not find or load main class Run.java
fix1: java Run

(a) First example.

cmd2: java Meta.java
err2: Could not find or load main class Meta.java
fix2: java Meta

(a) Second example.

match [STR(java), STR(Run.java)]
and [STR(Could), STR(not), STR(find), STR(or), STR(load), STR(main), STR(class), STR(Run.java)]
$\rightarrow_s$ [FSTR(java), FSTR(Run)]

(a) Symbolic rule representation synthesized after first example.

match [STR(java), VAR-MATCH$(1, \varepsilon,$ .java)]
and [STR(Could), STR(not), STR(find), STR(or), STR(load), STR(main), STR(class), VAR-MATCH$(2, \varepsilon,$ .java)]
$\rightarrow_s$ [FSTR(java), $\left\{ \begin{array}{l} \text{SUB-LR}(\text{IPOS}(0), \text{IPOS}(-5), \varepsilon, \varepsilon, \text{VAR}(1))]) \\ \text{SUB-LR}(\text{IPOS}(0), \text{CPOS}(.,1,0), \varepsilon, \varepsilon, \text{VAR}(1))]) \\ \text{SUB-LR}(\text{IPOS}(0), \text{CPOS}(.,-1,0), \varepsilon, \varepsilon, \text{VAR}(1))]) \end{array} \right\}$]

(a) Symbolic rule representation synthesized after both examples.

Figure 7: Two input examples e_1 and e_2 and symbolic rules synthesized after processing e_1 and e_2.

to unify the command **cmd2** in Figure 4a with the matching part of the already computed rule in Figure 5a.

rule:	STR(java)	STR(Run.java)
new-ex:	java	Meta.java
new-rule:	STR(java)	VAR-MATCH$(1, \varepsilon,$ .java)

2. All the fix expressions that cannot be constants are "promoted" to SUB-LR expressions that are consistent with the current examples and are allowed to use the variables appearing in the match expressions.

The second rule in Figure 7(d) reflects this update. The figure also shows how multiple SUB-LR expressions are represented symbolically as a set. We describe all of these components in detail in the next section.

4.3 Synthesis algorithm

Given a list of input examples, the function SYNTHRULES uses the first example and the function CONSTRULE to generate the symbolic rule composed only of constant operators. It then iteratively refines the rule on the remaining examples as shown in Figure 8. This second operation is done by the function REFINERULE which takes as input a symbolic rule r, one new example $(\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})$, and the list of examples E on which every concrete rule represented by r behaves correctly. REFINERULE executes two main steps using the following functions.

FINDVARIABLES tries to unify the inputs $\bar{s}_{cmd}$ and $\bar{s}_{err}$ with the corresponding match expressions cmd and err in the symbolic rule r and generates new variable match expressions if necessary—i.e., when r contains a matching expression STR(s) but the corresponding component in the example is a string different from s. In this case, a VAR-MATCH(i, l, r) expression is generated such that i is a new variable name, and l and r are the longest prefix and suffix shared by s, respectively.

When FINDVARIABLES is presented with a new $\bar{s}_{cmd}$ or $\bar{s}_{err}$ after a constant match expression has been 'promoted' to a VAR-MATCH(i, l, r) expression, the prefix and suffix are updated accordingly. FINDVARIABLES determines the longest prefix r' and suffix l' of l and r, respectively, that is consistent with the appropriate component of the new example, and generates VAR-MATCH(i, l', r').

SYNTHFIX uses the variables computed in the previous step to synthesize all possible fix expressions that are consistent with the list of examples $\{(\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})\} :: E$.

In order to simplify variable naming and guarantee unique names, each variable has the index of the corresponding element in the input—i.e., VAR(i) denotes the i-th string in the list $\bar{s}_{cmd}@\bar{s}_{err}$ obtained by concatenating the command and error input lists.

Lazy pattern matching The function FINDVARIABLES, given a rule r and a new example e, iterates over the input components of the new example e and outputs the set of variables necessary to match this new example with respect to the previously computed symbolic rule r. The function SYNTHFIX, given a rule r and a list of examples E, individually synthesizes all the components f_i of the symbolic output fix expression that are consistent with E. SYNTHFIX is incremental in the sense that it tries to minimally change the original fix expression of r:

- if the i-th component t_i of the fix expression of r is a constant string consistent with the new example, then it is left unchanged;
- in any other case the output has to be a substring operation, and the function SYNTHSUBSTRINGS is used to compute all the possible SUB-LR expressions that are consistent with the set of examples E.

Given a list of examples $e :: E$, the function SYNTHSUBSTRINGS first synthesizes all the SUB-LR expressions that are consistent with e using the function ALLSUBSTRINGS and then runs each synthesized expression on examples in E to remove the inconsistent ones.

ALLSUBSTRINGS Figure 8 omits the formal definition of the function ALLSUBSTRINGS due to space limitations, but we describe its main components. Given an example $e = (\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})$, a set of variable names V, and the index i corresponding to the element of the output sequence we are trying to synthesize, ALLSUBSTRINGS computes the set of all substring expressions of the form $fun =$ SUB-LR$(p_L, p_R, l, r,$ VAR$(j))$ that are consistent with e such that the result of applying fun to the j-th string in $\bar{s}_{cmd}@\bar{s}_{err}$ is the i-th string in $\bar{s}_{fix}$. Let's assume that $|\bar{s}_{cmd}| + |\bar{s}_{err}| = n_I$, $|\bar{s}_{fix}| = n_O$, and n_L is the length of the longest string appearing in any of the three lists in the input example. To compute the set ALLSUBSTRINGS(e, V, i) we iterate over all variable indices and for each variable index $j \in V$ we do the following.

1. Extract the string s_j corresponding to the variable VAR(j) — $O(n_I)$ iterations.
2. For each string s that is a substring of both $\bar{s}_{fix}[i]$ and s_j, compute all possible pairs of indices k_1, k_2 such that substr$(s_j, k_1, k_2) = s$ — $O(n_K^2)$ possible substrings and $O(n_K)$ possible ways to place the substring in $\bar{s}_{fix}[i]$.

```
//Rules consistent with input examples
fun SynthRules([e_0, ..., e_n])
    r ← ConstRule(e_0)
    for 1 ≤ i ≤ n do                          ▷ refine on each example e_i
        r ← RefineRule(r, [e_0, ..., e_{i-1}], e_i)
    return r

//Constant rule consistent with one example
fun ConstRule([s^c_1, ..., s^c_n], [s^e_1, ..., s^e_m], [s^f_1, ..., s^f_l])
    cst_cmd ← [Str(s^c_1), ..., Str(s^c_n)]
    cst_err ← [Str(s^e_1), ..., Str(s^e_m)]
    cst_fix ← [Fstr(s^f_1), ..., Fstr(s^f_l)]
    return (match cst_cmd and cst_err →_s cst_fix)

//Refines a rule to make it consistent with one more example
fun RefineRule(r, E, (s̄_cmd, s̄_err, s̄_fix))
    r ≡ (match cmd and err →_s fixes)
    (cmd', V_c) ← FindVariables(s̄_cmd, cmd, 0)
    (err', V_e) ← FindVariables(s̄_err, err, |s̄_cmd|)
    V ← V_c ∪ V_e
    E' ← (s̄_cmd, s̄_err, s̄_fix) :: E
    fixes' ← SynthFix(s̄_fix, fixes, E', V)
    return (match cmd' and err' →_s fixes')

//Finds variables necessary to match example
fun FindVariables([s_1, ..., s_n], [t_1, ..., t_m], o)
    if n ≠ m then                             ▷ Input length same as match length?
        return ⊥
    (m, V) ← ([], ∅)
    for 1 ≤ i ≤ n do
        case t_i = Str(s) ∧ s = s_i
            (m, V) ← (m@Str(s_i), V)
        case t_i = Str(s) ∧ s ≠ s_i
            pref ← LongestCommonPrefix(s, s_i)
            suf ← LongestCommonSuffix(s, s_i)
            newId ← i + o                     ▷ Create new variable
            m ← m@[Var-Match(newId, pref, suf)]
            V ← V ∪ {newId}
        case t_i = Var-Match(j, l, r)
            pref ← LongestCommonPrefix(s, l)
            suf ← LongestCommonSuffix(s, r)
            m ← m@[Var-Match(j, pref, suf)]
            V ← V ∪ {j}
    return (m, V)

//Outputs the fixes consistent with all the examples E and such that Sub-lr expressions can depend on any variable in V. The fix component of the latest example and the fixes computed on the previous examples are also passed as input
fun SynthFix([s_1, ..., s_n], [t_1, ..., t_n], e :: E, V)
    if n ≠ m then
        return ⊥
    f ← []
    for 1 ≤ i ≤ n do
        if t_i = Fstr(s) ∧ s = s_i then
            f ← f@[Fstr(s_i)]
        else
            f ← f@[SynthSubstrings(e :: E, V, i)]
    return f

//Outputs all Sub-lr expressions consistent with all the examples that can appear at position i in the fix expression
fun SynthSubstrings(e :: E, V, i)
    F ← AllSubstrings(e, V, i)
    for all (s̄_cmd, s̄_err, s̄_fix) ∈ E do
        F' ← ∅
        for all fun ∈ F do
            let Sub-lr(p_L, p_R, l, r, Var(j)) = fun
            o ← |s̄_cmd|
            if j < o then                     ▷ The variable is in s̄_cmd
                if Eval(fun, s̄_cmd[j] = s̄_fix) then
                    F' ← fun :: F'
            else                              ▷ The variable is in s̄_err
                if Eval(fun, s̄_err[j − o] = s̄_fix) then
                    F' ← fun :: F'
        F ← F'
    return F

//Outputs all Sub-lr expressions consistent with one example that can appear at position i in the fix expression
fun AllSubstrings((s̄_cmd, s̄_err, s̄_fix), V, i)
    return all Sub-lr(p_L, p_R, l, r, Var(j)) that
        when evaluated on s̄_cmd and s̄_err output s̄_fix[i]
        and such that j ∈ V.
```

Figure 8: Algorithm for synthesizing FIXIT rules from concrete examples.

3. For each k_1 (resp. k_2) compute every position expression p_1 (resp. p_2) such that evaluating p_1 (resp. p_2) on s_j produces the index k_1 (resp. k_2) — $O(n_K)$ possible positions.
4. For each of these possibilities yield the expression Sub-lr$(p_1, p_2, l, r, \text{Var}(j))$ where l and r are such that $\bar{s}_{fix}[i] = l \cdot \texttt{substr}(s_j, k_1, k_2) \cdot r$.

AllSubstrings produces a set of expressions that in the worst case has size $O(n_I n_K^5)$[3]. If we restrict the offset component δ to only range over the values $\{-1, 0, 1\}$ for the symbolic expressions Cpos(c, i, δ), the size reduces to $O(n_I n_K^3)$, and the synthesis algorithm is still sound and complete for this fragment of FIXIT.

This last restriction of the language can capture all the rules we are interested in. Notice that this analysis still holds in the extreme case in which all input matches are variable expressions of the form Var-Match$(i, \varepsilon, \varepsilon)$. In our experiments on real-world commands, worst-case performance is uncommon, and is induced by substring operations over heterogeneous strings which yield many possible implementations. Consider the following two examples.

cmd1: aaaa aaaa	**cmd2:** bbbb bbbb
err1: aaaa aaaa	**err2:** bbbb bbbb
fix1: aa	**fix1:** bb

Performing synthesis on this pair of examples yields a pattern match consisting of four Var-Match(i, ϵ, ϵ) expressions. Due to the uniformity of the input strings, synthesis yields 48 possible

[3] Note that the efficiency of this implementation is contingent on our specific choice of data structure and algorithms. A more naive solution, based on set intersection (like that of the DAG-based algorithm in FlashFill) may experience exponential blow-up in the number of examples, due to the quadratic nature of the intersection operations.

SUB-LR expressions. In particular, the desired fix can be generated from any of the four strings in the supplied s_{cmd} and s_{err}. Each of the four strings has three substrings of length 2, any of which yields the desired output. For each such substring, there are four pairs of IPOS values that supply the appropriate indices: The pair with two positive indices, the pair with two negative indices, and the two pairs consisting of one positive and one negative index.

Key point At this point we are ready to explain why all the match expressions can be kept as constants for as long as possible. If after processing a set of examples E, some expression in cmd or err is of the form $\text{STR}(s_i)$, then, for every input example, the value of the i-th component is the string s_i. Therefore, even if we replace this expression with a variable, all its instantiations will have the same values. Consequently, every function of the form $\text{SUB-LR}(p_1, p_2, l, r, \text{VAR}(i))$ will produce a constant output, making it equivalent to the some constant function $\text{FSTR}(s')$.

4.4 Partition-Based Synthesis

FIXIT can learn repair rules from a set of examples corresponding to a specific incorrect use of a command. In practice, however, it may be difficult to present FIXIT with a collection of neatly curated sets of examples, from each of which, FIXIT learns a single symbolic rule. Such a process is both labor-intensive and error-prone. We instead envision large-scale learning of FIXIT rules from undifferentiated sets of examples submitted by many users. To facilitate this, we propose a simple multi-stage partitioning strategy. As a consequence of FIXIT's structure, each symbolic rule matches on a pair of command and error strings, each of which has a fixed number of tokens. Upon a match, FIXIT generates a repaired command with a fixed number of tokens. Conversely, every symbolic FIXIT expression must be learned from a set S_E of example triples of the form $(s_{cmd}, s_{err}, s_{fix})$ for which the lengths of s_{cmd}, s_{err}, and s_{fix} do not vary.

Given an undifferentiated set of examples, S_E, we partition S_E into n disjoint subsets δ_i where $S_E = \bigcup_{i=0}^{n} \delta_i$. For every δ_i the property $\forall (s_{cmd}, s_{err}, s_{fix}) \in \delta_i,\ (\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix}) \in \delta_i . |s_{cmd}| = |\bar{s}_{cmd}| \wedge |s_{err}| = |\bar{s}_{err}| \wedge |s_{fix}| = |\bar{s}_{fix}|$. This divides S_E into subsets from which it is possible to synthesize FIXIT rules.

After dividing S_E, it may still be the case that individual sets δ_i contain examples representing distinct command repair rules which share the same triple of s_{cmd}, s_{err}, and s_{fix} lengths. At this point, we attempt to find the smallest set of rules that can be synthesized from the examples in δ_i.

The search ranges over all partitions P of δ_i, where P is the set $\{\delta_{i,1}, \delta_{i,2}, \ldots, \delta_{i,m}\}$ such that $\bigcup_{j=1}^{m} \delta_{i,j} = \delta_i$ and $\forall j \neq k, \delta_{i,j} \cap \delta_{i,k} = \varnothing$. We enumerate the partitions in ascending order of size m, starting with $P = \{\delta_i\}$, and ending with the partition consisting entirely of singleton sets. Given a partition P, we attempt to synthesize a symbolic FIXIT rule for each set in P, stopping when we have generated a rule for each $\delta_{i,j}$. As we will show in section 6.2, this expensive procedure is only feasible when using our novel contribution of lazy VSA.

Example partitioning Consider the example set shown in Figure 9. The first two examples in this set correspond to the repair in Section 2.1 while the last two correspond to the repair in Section 2.3. The third and fourth example correspond to the rule that outputs `composer`, followed by the token at index 8 in the input (i.e., the first token suggested by the error message).

We first group the examples based on the length of their components. For the first four examples, we have $|s_{cmd}| = 2$, $|s_{err}| = 8$, $|s_{fix}| = 2$ and for the two remaining examples, $|s_{cmd}| = 3$, $|s_{err}| = 8$, $|s_{fix}| = 6$. Thus, we obtain two groups $S_1 = \{e_1, e_2, e_3, e_4\}$ and $S_2 = \{e_5, e_6\}$.

e_1	**cmd1:**	java Run.java
	err1:	Could not find or load main class Run.java
	fix1:	java Run
e_2	**cmd2:**	java Test.java
	err2:	Could not find or load main class Test.java
	fix2:	java Test
e_3	**cmd3:**	composer pkg
	err3:	did you mean one of these? pkg1 pkg2
	fix3:	composer pkg1
e_4	**cmd4:**	composer hptt
	err4:	did you mean one of these? http html
	fix4:	composer http
e_5	**cmd5:**	mv photo.jpg Mary/summer12.jpg
	err5:	can't rename 'photo.jpg': No such file or directory
	fix5:	mkdir Mary && mv photo.jpg Mary/summer12.jpg
e_6	**cmd6:**	mv dec31.jpg Bob/family.jpg
	err6:	can't rename 'dec31.jpg': No such file or directory
	fix6:	mkdir Bob && mv dec31.jpg Bob/family.jpg

Figure 9: Examples requiring more than one FIXIT rule.

While there exists a FIXIT program that is consistent with the examples in the set S_2, no FIXIT program can describe a transformation that is consistent with all the examples in S_1. We therefore proceed by iteratively partitioning the set S_i, attempting to find a partition P_i for which we can synthesize a rule for every $S_{i,j} \in P_i$. For S_2, we can clearly do so for the initial partition $P_2 = \{S_2\}$, yielding the rule from Section 2.2. For S_1, the first partition for which we can generate a FIXIT rule for every element is $P_1 = \{S_{1,1} = \{e_1, e_2\},\ S_{1,2} = \{e_3, e_4\}\}$. S_1 yields the rule from Section 2.1, and S_2 yields the simple substitution rule described previously.

4.5 More Succinct Representation

The version of symbolic rule we presented is already able to store exponentially many concrete FIXIT rules in polynomial space. In this section, we discuss further improvements that can make the representation more succinct.

Avoid redundancy In the set of fix expressions enumerated by the function ALLSUBSTRINGS, the last three components of the the expression $\text{SUB-LR}(p_1, p_2, l, r, \text{VAR}(j))$ are often repeated many times. Looking at Figure 6a we can see how all the synthesized functions have $l = r = \varepsilon$ and are applied to the variable $\text{VAR}(1)$. We define a data structure for representing sets of fix expressions that avoids these repetitions. A set of fix expressions is represented symbolically using a partial function

$$d : \mathbb{N} \mapsto (\Sigma^* \times \Sigma^*) \mapsto Set(P \times P)$$

where P is the set of all position expressions. Formally, given a variable index i and two strings l and r, the set $d(i, l, r)$ symbolically represents the set of fix expressions $\{\text{SUB-LR}(p_1, p_2, l, r, \text{VAR}(i)) \mid (p_1, p_2) \in d(i, l, r)\}$. The function d can be efficiently implemented and avoids redundancy. Considering again the example rule in Figure 6a, all the fix expressions in the second component of the output can be succinctly represented by the function d that is only defined on the input $(1, \varepsilon, \varepsilon)$ and such that

$$d(1, \varepsilon, \varepsilon) = \{\ \begin{array}{l}(\text{IPOS}(0), \text{IPOS}(-5)),\\ (\text{IPOS}(0), \text{CPOS}(.\,, 1, 0)),\\ (\text{IPOS}(0), \text{CPOS}(.\,, -1, 0))\end{array}\ \}.$$

Avoid example representation Each synthesized rule is in some sense coupled to the set of examples used to synthesize it. We present a data structure that only keeps track of the important

"parts" of the input examples and therefore allows us to discard each example after it has been processed.

We modify the symbolic rule representation as follows. Given a set of examples $(\bar{s}^1_{cmd}, \bar{s}^1_{err}, \bar{s}^1_{fix}) \dots (\bar{s}^n_{cmd}, \bar{s}^n_{err}, \bar{s}^n_{fix})$,

- every variable VAR-MATCH(i, l, r) in the match component becomes a pair (VAR-MATCH$(i, l, r), [b_1, \dots, b_n]$) where the second component is the list of strings that binds to VAR(i) in the input components of the examples—i.e. $b_j = (\bar{s}^j_{cmd} \bar{s}^j_{err})[i]$;
- every set of fix expressions represented by a function d_i and corresponding to the i-th component of the fix expression becomes a pair $(d_i, [b_1, \dots, b_n])$ where the second component is the list of strings that appear in position i in the output components of the examples—i.e. $b_j = \bar{s}^j_{fix}[i]$;

Using this data structure we do not need to store examples as we can always re-infer them from the symbolic rule representation.

4.6 Concrete outputs

Taking into account the updated data structures, the algorithm SYNTHRULES returns a symbolic rule r of the form **match** cmd **and** $err \ \rightarrow_s \ fixes$ where $cmd = [c_1, \dots, c_n]$ and $err = [e_1, \dots, e_m]$ are lists of expressions of the form STR(s) or (VAR$(i), B$), while $fixes = [f_1, \dots, f_l]$ is a list of expressions of the form FSTR(s) or (d, B) where d is the data structure for representing multiple fix expressions. The set of concrete FIXIT rules induced by this symbolic representation is the following.

$$\begin{aligned}
&\texttt{con}(\textbf{match}\ cmd\ \textbf{and}\ err\ \rightarrow_s\ fix) = \\
&\qquad \{\textbf{match}\ cmd\ \textbf{and}\ err\ \rightarrow\ f \mid f \in \texttt{con}(fix)\} \\
&\texttt{con}([f_1, \dots, f_l]) = \{[f'_1, \dots, f'_l] \mid f'_i \in \texttt{con}(f_i)\} \\
&\texttt{con}(\text{FSTR}(s)) = \{\text{FSTR}(s)\} \\
&\texttt{con}(d, B) = \{\text{SUB-LR}(p_1, p_2, l, r, \text{VAR}(i)) \mid \\
&\qquad\qquad \exists i, l, r.(p_1, p_2) \in d(i, l, r)\}.
\end{aligned}$$

5. Formal properties

We study the formal properties of the synthesis algorithm and of the language FIXIT. These specific properties describe the behavior of the synthesis algorithm in the absence of the partitioning strategy described in Section 4.4.

Properties of the synthesis algorithm First, our synthesis algorithm is invariant with respect to the order in which the training examples are presented. Thus, the properties of a symbolic rule generated by SYNTHRULES, can be discussed solely in terms of the *set* of examples provided to SYNTHRULES.

Theorem 1 (Order invariance)**.** *Given a list of examples E, for every permutation of examples E' of E, we have* $\texttt{con}(\text{SynthRules}(E)) = \texttt{con}(\text{SynthRules}(E'))$.

Proof sketch. Consider a list of examples E. If two examples differ at the ith position in their respective commands, the ith expression in cmd will be promoted to a VAR-MATCH, regardless of the order in which they are presented to SYNTHRULES. Moreover, if all $|E|$ strings at the ith position in the commands share a prefix or suffix, reordering E does not change this fact. Thus, the discovered VAR-MATCH expressions will not vary based on order. The same holds for err and the "promotion" of constants in SYNTHFIX.

Moreover, SYNTHSUBSTRING starts from scratch at each iteration of the loop in SYNTHRULES and fix depends only on the output of SYNTHFIX in the final iteration. Since the variable set does not vary based on the ordering of E, the final invocation of SYNTHFIX does not depend on the ordering of E. □

Second, the synthesis algorithm produces only rules that are consistent with the input examples. If we select an arbitrary concrete rule r from the set specified by a symbolic rule generated by SYNTHRULES, and run it on the command and error of any of the examples provided to SYNTHRULES for the synthesis of r, we will obtain the fix originally provided in that example.

Theorem 2 (Soundness)**.** *Given a list of examples E, for every rule* $r \in \texttt{con}(\text{SynthRules}(E))$ *and for every example* $(\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix}) \in E$, $[\![r]\!](\bar{s}_{cmd}, \bar{s}_{err}) = \bar{s}_{fix}$.

Proof sketch. The repeated applications of FINDVARIABLES will promote any STR(s) expression if a new example does not match on s. Moreover, when refining a VAR-MATCH, FINDVARIABLES chooses the longest prefix and suffix consistent with all examples seen so far. Thus, cmd and err will correctly match on all examples. The soundness of the resulting fix derives from the fact that at each iteration of the loop in SYNTHRULES, the invocation of SYNTHSUBSTRING in SYNTHFIX takes into account all examples seen in previous iterations of the loop. Moreover, each invocation begins with the set of all possible SUB-LR expressions, and prunes those inconsistent with any example seen so far. □

Since parts of the match expressions are "promoted" to variables only when the input examples show that this is required, our synthesis algorithm does not explicitly keep track of all the possible rules that can be consistent with the examples. Our completeness result reflects this idea.

Theorem 3 (Completeness)**.** *Given a set of examples E, for every* Fixit *rule r that is consistent with E, either* $r \in \texttt{con}(\text{SynthRules}(E))$ *or there exists an example e such that* $r \in \texttt{con}(\text{SynthRules}(e :: E))$.

Proof sketch. Concretely, a particular rule r that is consistent with E might not appear in $R = \text{SYNTHRULES}(E)$. However, this can only happen because the match expression of r has more variables than the match of any rule in R. This can be fixed by providing an example that forces the algorithm to promote to variables all the required match expressions. □

***Properties of the language* Fixit** We define the size of an input example $e = (\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})$ and the size of each rule r consistent with e as the sum of its lengths $\texttt{size}(e) = \texttt{size}(r) = |\bar{s}_{cmd}| + |\bar{s}_{err}| + |\bar{s}_{fix}|$. For each set of examples, there can be exponentially many rules consistent with it.

Theorem 4 (Number of consistent rules)**.** *Given a set of examples E, each of size k, the set* $\texttt{con}(\text{SynthRules}(E))$ *contains* $2^{O(poly(k))}$ *rules.*

Proof sketch. As we showed in Section 4.3, for each position in the output of a FIXIT rule there are potentially polynomially many SUB-LR expressions consistent with the provided examples. For the ith position in fix, we are free to choose any of the possible SUB-LRs, independently of our choice at other positions. Thus, the number of possible rules is potentially exponential in $|\bar{s}_{fix}|$. □

Despite the exponential number of rules represented, our data structure allows the SYNTHRULES to encode these rules using only polynomial size.

Next, there exists an active learning algorithm for learning FIXIT rules that requires only a polynomial number of examples—i.e., queries to the user.

Theorem 5 (Complexity of active learning)**.** *If there exists a target rule r of size k, there exists an active learning algorithm that will learn r by asking $O(poly(k))$ queries of the form:* What should the output of r be on the input $(\bar{s}_{cmd}, \bar{s}_{fix})$.

Proof sketch. The algorithm first asks k queries to figure out which match expressions are variables and which ones are constants. Then, for each output component for which there exists two possible fix expressions SUB-LR consistent with the examples, it asks

a query that differentiates the two. Since there are only $O(kn_I n_K^5)$ many expressions in the output the algorithm will ask at most polynomially many queries. □

6. Implementation and Evaluation

We now describe the implementation details of NOFAQ, as well as our experimental evaluation of NOFAQ on a set of examples and test cases isolated from THEFXXX and web forums.

6.1 Implementation

We implemented the language FIXIT and its synthesis algorithm in a system called NOFAQ. NOFAQ is implemented in F#[4] and consists of some additional optimizations and design choices as described below.

6.1.1 Implementation optimizations

The function ALLSUBSTRINGS in Figure 8 synthesizes all SUB-LR functions that are consistent with the first input/output pair (s, t) of strings in the example set E and then applies each of the synthesized functions to the other elements of E for filtering only consistent functions. In practice, we first compute the longest common prefixes and suffixes of the strings appearing in the components $\bar{s}_{fix}$ of E to avoid enumerating instances of the form SUB-LR$(_, _, l, r, _)$ such that l or r are not prefixes or suffixes of some output string t appearing in E.

The other optimization is based on the following property of the REFINERULE function: when adding a new example to r, if the function FINDVARIABLES introduces a new set of variables V, all the new instances of SUB-LR that did not already appear in r depend on one of the newly introduced variables in V. Based on this idea, the function ALLSUBSTRINGS only has to compute functions of the form SUB-LR$(_, _, _, _,$ VAR$(i))$ where $i \in V$, and can reuse the previously computed functions for the other variables by simply filtering the ones that behave correctly on the newly introduced example.

6.1.2 Ranking

Since there can be multiple possible expressions in FIXIT that are consistent with the examples, we employ a simple ranking technique to select an expression amongst them. If there are multiple SUB-LR expressions that can generate the desired output string, we select the expression that uses the variable with the lowest index—i.e., the leftmost one. Similarly, the l and r included in VAR-MATCH expressions implicitly encode all rules matching on prefixes and suffixes of l and r, respectively. We choose the longest l and r over all others.

As the example set increases in size, we envision users will likely submit diverse sets of examples, particularly in use cases with thousands of users submitting examples. As users submit examples which draw from heterogeneous collections of command parameters, VAR-MATCH prefixes and suffixes should converge to the least restrictive versions. Similarly NOFAQ should discover the least restrictive set of constants for both match expressions. As these input parameters vary over the set of examples, spurious ambiguities in SUB-LR should be eliminated when NOFAQ is presented with specific fix examples which function as counter-examples to unnecessary substring expressions.

6.2 Evaluation

We now describe our experimental evaluation. The experiments were run on an Intel Core i7 2.30GHz CPU with 16 GB of RAM. We present both qualitative and quantitative analysis of the algorithm. We assess the expressiveness of NOFAQ by attempting synthesis on a benchmark suite that includes the rules in the tool THEFXXX. We then evaluate the performance of NOFAQ and its scalability.

6.2.1 Benchmark Suite

We compiled our benchmark suite from an initial set of of 92 benchmarks, which were collected from both THEFXXX (76) and online help forums (16). We considered the 76 repair rules hard-coded in the THEFXXX tool to assess the expressiveness of NOFAQ. Since rules in THEFXXX can use arbitrary Python code, it is hard to exactly compare them to the ones produced by NOFAQ. We use manual testing to check that a rule r generated by our tool is *consistent* with a rule r' in THEFXXX. To do so, we manually constructed a set of examples based on the pattern-matching and textual substitutions performed by the THEFXXX rules.

The other sixteen example sets were obtained from examples found during a non-exhaustive survey of command-line help forums on the web. These commands consist of various types of git, svn, and mvn commands, including committing, reverting, and deleting from repositories, as well as installing and removing packages.

The NOFAQ system is able to synthesize a rule for 81 of the 92 benchmarks. The remaining 11 failing benchmarks can be divided into three broad categories: i) Hard-coded operations searching for specific strings in some context (8), ii) Complex patterns checking relationships between variable expressions (2), and iii) Error messages displaying parts of the input file's content (1). We did not provide examples for these 11 rules. We elaborate more about these categories in Section 7.

Number of examples We observed that it was natural to provide two to five examples per benchmark for NOFAQ to uniquely learn the desired fix rule. We also provided additional examples for manually testing the learned rules, yielding a set of three to six examples. Given the rule appearing in Figure 7, for example, we used the two examples in Figure 7 and another example with the file name Employee.java. In future, we envision users to contribute different examples to the system for automatically building a large corpus of learned fix rules.

While these examples are synthetic examples reverse engineered from other sources, they are also natural examples which exercise the range of e.g. path and file names one would expect to see in a real Unix system. In the case of the repaired command in Section 2.3, the natural two-example set would consist of two distinct directory names which do not share prefixes and suffixes, as well as filenames with distinct prefixes and extensions.

6.2.2 Qualitative evaluation

Given a single set containing examples for all the 81 cases in which NOFAQ is capable of synthesizing a rule, we performed synthesis as described in Section 4.4. For each rule we retained a single example from the training set and used it to test the accuracy of each rule. We also report how often a given input could be repaired using more than one rule.

Results For all 81 cases, NOFAQ synthesized a rule consistent with the corresponding THEFXXX rule or web forum answer. In some cases we had to synthesize more than one FIXIT rule to capture the different possible behaviors of a single rule in THEFXXX. For example, one can try adding 'sudo' in front of a command for several possible errors such as "Command not found", "You don't have the permission" etc. In such cases, thanks to the partitioning algorithm, NOFAQ generated a separate rule for each possible error message. For each case where we synthesized a rule, correctness

[4] The implementation will be made open-source and publicly available after the review process.

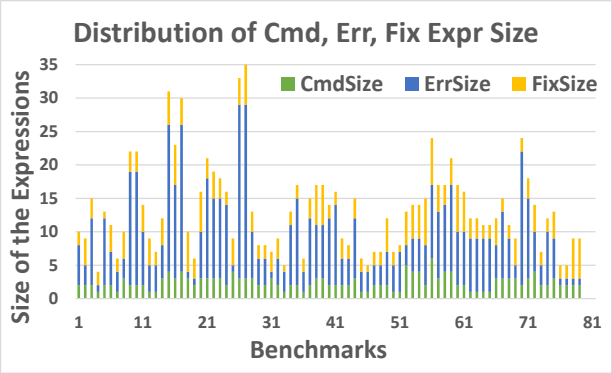

Figure 10: The distribution of benchmarks in terms of individual sizes of $\bar{s}_{cmd}$, $\bar{s}_{err}$, and $\bar{s}_{fix}$ expressions in the examples.

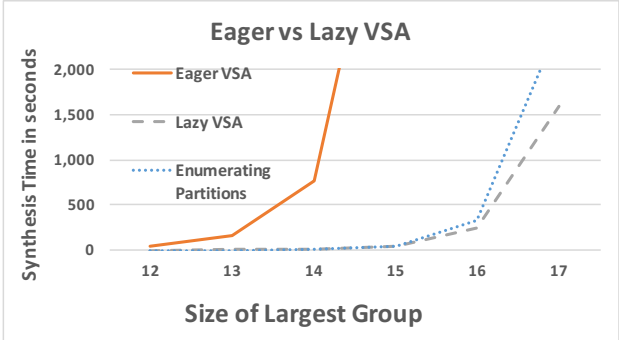

Figure 11: Synthesis times for different benchmarks for the lazy and non-lazy rule representations.

was independent of our choice of examples. If the correct rule was synthesized, it was synthesized regardless of which subset of the examples provided for that rule we selected.

Distribution of rule sizes We define the size of an expression such as $\bar{s}_{cmd}$, $\bar{s}_{err}$, and $\bar{s}_{fix}$ as the number of strings present in it. The distribution of the size of the benchmarks in terms of the sizes of the $\bar{s}_{cmd}$, $\bar{s}_{err}$, and $\bar{s}_{fix}$ tuples in input-output examples is shown in Figure 10. Note that we do not show two benchmarks in the graph with disproportionately high $\bar{s}_{err}$ expression size of 110 for clarity. The average total size of the examples in the benchmarks was 15.91 ± 17.18[5], with the maximum size of 116. The average sizes for the individual expressions of the examples were: i) $\bar{s}_{cmd}$: 2.38 ± 1.01 with maximum of 6, ii) $\bar{s}_{err}$: 10.12 ± 16.85 with a maximum of 110, and iii) $\bar{s}_{fix}$: 3.41 ± 1.55 with a maximum of 7.

Distribution of rule matching For each set of example provided for an individual rule, we isolated one example to measure the accuracy of the tool. All the test examples were correctly described by at least one of the synthesized rules. For the majority of the test cases, there was exactly one rule which matched both the command and error message. The remaining 12 test cases which matched against multiple rules came from collections of example sets which represented different fixes of the same command and error messages.

Total test cases	81
Test cases matched by one rule	69
Test cases matched by multiple rules	12

Ranking Consistent with our hypothesis in Section 6.1.2, a diverse set of examples was sufficient for eliminating spurious restrictions and substring expressions. In every test case, the rule chosen by our ranking policy was capable of correcting all test cases presented. In practice, many rules still have several possible correct SUB-LR expressions. However, this remaining ambiguity occurs because the same string can appear many times in the command and error message (e.g., the string Employee in the example in Section 2.1).

6.2.3 Quantitative evaluation

We now report on the quantitative metrics of our synthesis algorithm. In this section we only report data for the 81 benchmarks for which NOFAQ can successfully synthesize a FIXIT rule.

[5] We use $a \pm s$ to denote an average a with standard deviation s.

Evaluation of lazy VSA synthesis time In Figure 11, we show the time taken to partition and synthesize FIXIT rules for the 81 benchmarks, using both the lazy and a non-lazy rule representation, as the number of examples per benchmark increases. The non-lazy representation always considers match and fix expressions as variables, rather than initially starting with constants.

To test the performance of the lazy and non-lazy representations as the size of the input set increases, we iteratively increase the size of the training set. For each test, we add a single example to one of the benchmarks and then attempt synthesis. We plot the synthesis time with respect to the largest set of examples for which we must enumerate possible partitions until we successfully synthesize rules. To understand the performance overhead induced by synthesis, we also evaluate a version of the algorithm which enumerates partitions without performing synthesis. For each algorithm, we iteratively increased the training set size until the algorithm reached a $2,000$ second timeout.

The non-lazy VSA incurs a significant overhead, and scales much worse than the lazy version, reaching the timeout when the largest set has 14 examples. The lazy VSA, in contrast, is much closer to the optimum; the synthesis time is negligible compared to the inherent cost of enumerating all partitions of a set. In fact, the lazy synthesis actually completes faster than exhaustive enumeration. This is reasonable, as the first partitioning which yields a successful FIXIT rule for all subsets tends to be somewhere near the middle of the enumeration, and thus does not incur the cost of enumerating the remainder of the search space. In summary, the lazy VSA strictly outperforms non-lazy VSA and can handle much larger sets of examples.

SUB-LR *expression in synthesized rules* The distribution of FSTR and SUB-LR expressions in the synthesized FIXIT rules is shown in Figure 12. The output components of the synthesized rule contain on average $29.01\% \pm 24.4\%$ SUB-LR expressions. Concretely, a synthesized rule contains on average 0.91 ± 0.76 SUB-LR expressions.

Synthesis time vs. number of* SUB-LR *expressions The synthesis time for different numbers of SUB-LR expressions in the repair rule is shown in Figure 13. As expected, the benchmarks that do not contain SUB-LR expressions take negligible time. The benchmarks involving two SUB-LR expressions on average require more time than the benchmarks with a single SUB-LR expression. Interestingly, the benchmarks with 3 SUB-LR expressions take lesser time than the benchmarks with 2 SUB-LR expressions. A possible explanation for this behavior is that the complexity of substring extraction tasks for these benchmarks is relatively simpler (e.g. identity) than the benchmarks with 2 SUB-LR expressions.

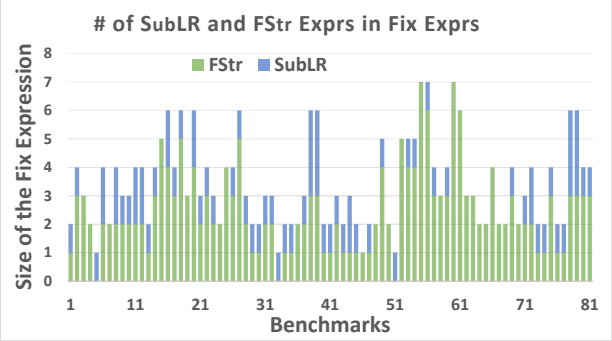

Figure 12: The distribution of FSTR and SUB-LR expressions in the final synthesized repair expression.

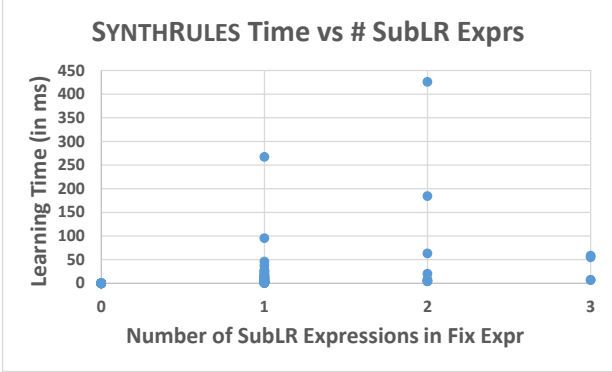

Figure 13: Synthesis times for varying number of SUB-LR expressions in the repair rule.

Scalability of synthesis algorithm with example size Since all real-world examples we collected are relatively of small size (with maximum size of 116 space-separated strings), we evaluate the scalability of the SYNTHRULES algorithm by creating artificial examples of increasing sizes. We create these artificial examples by repeating the $\bar{s}_{cmd}$, $\bar{s}_{err}$, and $\bar{s}_{fix}$ commands multiple times for the example shown in Section 2.2. The synthesis times for increasing size of examples is shown in Figure 14. From the graph, we observe that the synthesis times scale in a quadratic fashion with respect to the example size.

7. Limitations

We showed that the language FIXIT is able to express many real-world command line repair rules and that these rules can be synthesized using few examples. We now present some limitations of our approach, in particular with respect to the 11 THEFXXX rules that FIXIT could not describe.

Complex patterns Two rules were checking complex properties of the input that FIXIT cannot capture. For example, FIXIT cannot check whether the error message contains some special character. FIXIT's conditional matching is limited to whole string or prefix/suffix matching, and thus cannot check if e.g. a file name contains a non-unicode whitespace character. All character relative logic occurs in the substring generation after input matching. FIXIT also cannot check whether some string in the input command is repeated more than once.

Context-dependencies Eight rules had hard-coded operations that were searching some context (the file system, a configuration file, etc.) for specific strings to complete the output. FIXIT only receives as inputs the command and the error message, and the rules currently cannot use any context.

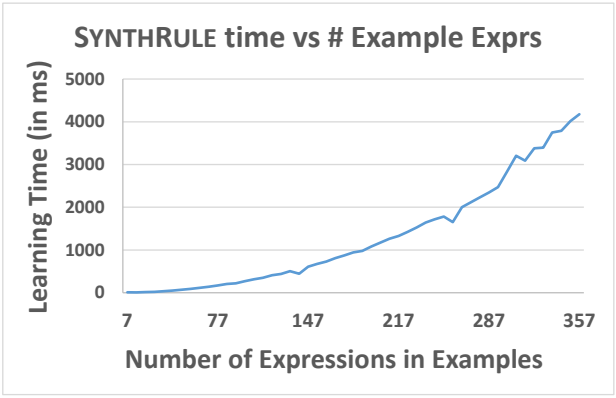

Figure 14: Synthesis times with increasing size of examples.

8. Related work

Version-space algebra for synthesis The concept of Version-space algebra (VSA) was first introduced by Mitchell [17] in the context of machine learning and was later used by Lau et al. to learn programs from demonstrations in a Programming By Examples/Demonstrations system called SmartEdit [11]. It has since been used for many PBE systems from various domains including syntactic string transformations in FlashFill [7], table transformations [9, 25], number transformations [24], text extraction from semi-structured text files in FlashExtract [12], and transformation of semi-structured spreadsheets to relational tables in FlashRelate [2]. Our synthesis algorithm also uses VSA to succinctly represent a large set of conforming expressions. However, in contrast to previous approaches that represent all conforming expressions concretely and then use intersection for refinement, our synthesis algorithm maintains a lazy representation of rules and concretizes the choices on demand in a lazy fashion only when it is needed. Moreover, our careful design of DSL operators and the corresponding VSA in NOFAQ lead to a polynomial time synthesis algorithm unlike most previous approaches that have exponential time synthesis algorithms.

In particular, it is illustrative to compare the FlashFill DSL with FIXIT. While, like FIXIT, FlashFill synthesizes string manipulations from input-output examples, specific performance properties make it less suitable for large scale learning from large sample sets. Prior to developing NOFAQ, we evaluated the possibility of simply using the FlashFill algorithm as-is for our purpose of learning command repair rules. Early empirical results indicated that the off-the-shelf algorithm scaled poorly as the error messages increased in length, which was a common occurrence for our benchmarks. Moreover, other limitations of no offset operator in position expressions and support for finite hard-coded regular expression tokens made FlashFill unsuitable for learning SUB-LR expressions.

We isolated several theoretical properties of FlashFill's key operators which yielded poor performance on large inputs. In particular, the binary concatenation operator over arbitrary substrings of

the entire input string induces a DAG structure for the symbolic representation of programs. More explicitly, given an example output string S, there exists a node n_p for each position p in S. An edge from n_p to n'_p, $p < p'$ represents the substring $S[p : p']$. Each edge is labelled with the set F of functions over the example inputs which yield the substring. Thus, a path from n_0 to $n_{|S|}$ represents some concatenation of the output of several string operations which yields the desired output. Given a DAG D consistent with a set of examples E, FlashFill incorporates a new example e represented by DAG D' by taking the cartesian product of the vertices of D and D', to construct a new DAG D''. An edge with label set F'' in the new DAG represents a set of functions which were part of a correct program for the examples E, and also map from the inputs of e to a substring of e's output. The iterated cartesian products yield time complexity exponential in the number of examples.

FIXIT, in contrast, posesses unary string operations constrained to specific variable terms. FIXIT's unary SUB-LR operator yields a language that is disjoint from FlashFill with concatenation removed; we obtain a language expressive enough for a large set of practical command repair transformations isolated from real use cases, while dramatically improving worst-case performance. The constrained nature of the program representation lets the synthesis algorithm eliminate programs inconsistent with a new example without directly computing the intersection of two sets of candidate programs, ensuring polynomial-time performance even in the worst case. Moreover, FIXIT also allows for repair transformations that require arbitrary offsets from a regex match, which are not expressible in the FlashFill DSL.

Programming by Examples (PBE) PBE has been an active research area in the AI and HCI communities from a long time [15]. In addition to VSA-based data wrangling [8], PBE techniques have recently been developed for various domains including interactive synthesis of parsers [14], synthesis of recursive functional programs over algebraic data types [4, 18], synthesizing sequence of program refactorings [20], imperative data structure manipulations [26], and network policies [31]. Our technique also learns repair rules from few input-output examples of buggy and fixed commands, but both our problem domain of learning command repairs and the learning techniques of using lazy VSA are quite different from these PBE systems.

Program repair Research in automated program repair focuses on automatically changing incorrect programs to make them meet a desired specification [6]. The main challenge is to efficiently search the space of all programs to find one that behaves correctly. The most prominent search techniques are enumerative or data-driven. GenProg uses genetic programming to repeatedly alter the incorrect program in the hope to make it correct [13]. Data-driven approaches use the large amount of code that is publicly available online to synthesize likely changes to the input program [21, 30]. Prophet [16] is a patch generation system that learns a probabilistic application-independent model of correct code from a set of successful human patches. Unlike these techniques that learn a global model of code repair across different applications, our technique learns command-specific repairs by observing how expert users fix their buggy commands — i.e., from both the incorrect command the user started with (together with the error message) and the correct command she wrote as a fix.

Crowdsourced Repairs *HelpMeOut* is a social recommender system that helps novice users facing programming errors by showing them examples of how other programmers have corrected similar errors [10]. While *HelpMeOut* can show examples of similar fixes it does not concretely show the user how the code should be corrected. This aspect is the major difference between *HelpMeOut* and NOFAQ.

THEFXXX provides a Python interface for command substitution and repair rules, and it requires a degree of language and tool-specific knowledge that may not be accessible to command line novices, particularly if non-trivial substring operations are required to derive the desired command. Much like FlashFill, we aim to emulate the workflow of non-technical users communicating with experts on web forums. For a beginner learning the command line, Python string manipulations are likely a fairly challenging task, and the cost of an incorrectly transformed shell command is potentially catastrophic. In such a situation where a non-expert desires a new THEFXXX rule, such a user may provide an example of several command/error pairs, and the desired fix for each, from which an expert would write the desired Python code. NoFAQ shortens this loop by moving the fix synthesis into a polynomial time algorithm on the user's machine.

Rule learning Rules provide a simple way to represent programmers actions and in general any type of data transformation. Rule learning has been extensively investigated in classical machine learning and data mining [5]. The goal of rule learning is to discover and mine rules describing interesting relations appearing in data. Common concept classes for describing rules are Horn clauses or association rules [19]. The approach presented in this paper differs from rule learning in two aspects: 1) the rules are expressed in a complex concept class and are hard to learn — i.e., FIXIT programs; 2) the examples are given by a teacher that has in mind a target rule. In the future we plan to build a system that uses rule learning techniques to mine FIXIT rules from unsupervised data.

Program synthesis There has been a resurgence in Program Synthesis research in recent years [1]. In addition to examples (as described above), there have been several techniques developed for handling other forms of specifications such as partial programs [27, 28], reference implementations [23], and concrete traces [29]. While these specification mechanisms have been found to be useful in several domains, we believe examples are the most natural mechanism for specifying command line repairs especially for beginners. There is also a recent movement towards using data-driven techniques for synthesis [22], e.g. the PLINY project (`http://pliny.rice.edu/index.html`). In future, we envision our system to also make use of large number of examples of buggy commands and their corresponding repairs to learn a big database of FIXIT rules.

9. Conclusion and Future Directions

We presented a tool NOFAQ that suggests possible fixes to common buggy commands by learning from examples of how experts fix such issues. Our language design walks a fine line between expressivity and performance: by careful choice of unary operators over pre-defined variables, and exclusion of arbitrary substring operations, we avoid exponential-time worst case performance, while still maintaining a useful degree of functionality. NOFAQ was able to instantly synthesize 85% of the rules appearing in the popular repair tool THEFXXX and 16 other rules from online help forums. Although NOFAQ tool is aimed towards repairing commands, we believe our novel combination of synthesis and rule-based program repair is quite general and is applicable in many other domains as well. We plan to to apply this methodology to more complex tasks, such as correcting syntax errors in source code, applying code optimization, and editing configuration files. In the future, we hope to create a tool which can take large command histories from expert users and quickly derive rules, as well as synthesize new rules online as experts use the shell.

References

[1] R. Alur, R. Bodík, E. Dallal, D. Fisman, P. Garg, G. Juniwal, H. Kress-Gazit, P. Madhusudan, M. M. K. Martin, M. Raghothaman, S. Saha, S. A. Seshia, R. Singh, A. Solar-Lezama, E. Torlak, and A. Udupa. Syntax-guided synthesis. In *Dependable Software Systems Engineering*, pages 1–25. 2015.

[2] D. W. Barowy, S. Gulwani, T. Hart, and B. Zorn. Flashrelate: Extracting relational data from semi-structured spreadsheets using examples. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2015, pages 218–228, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3468-6. doi: 10.1145/2737924.2737952. URL http://doi.acm.org/10.1145/2737924.2737952.

[3] S. M. Doane, D. S. McNamara, W. Kintsch, P. G. Polson, and D. M. Clawson. Prompt comprehension in unix command production. *Memory & cognition*, 20(4):327–343, 1992.

[4] J. K. Feser, S. Chaudhuri, and I. Dillig. Synthesizing data structure transformations from input-output examples. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation, Portland, OR, USA, June 15-17, 2015*, pages 229–239, 2015. doi: 10.1145/2737924.2737977. URL http://doi.acm.org/10.1145/2737924.2737977.

[5] J. Furnkranz, D. Gamberger, and N. Lavrac. Rule learning in a nutshell. In *Foundations of Rule Learning*, Cognitive Technologies, pages 19–55. Springer Berlin Heidelberg, 2012. ISBN 978-3-540-75196-0. doi: 10.1007/978-3-540-75197-7_2. URL http://dx.doi.org/10.1007/978-3-540-75197-7_2.

[6] C. Goues, S. Forrest, and W. Weimer. Current challenges in automatic software repair. *Software Quality Journal*, 21(3):421–443, Sept. 2013. ISSN 0963-9314. doi: 10.1007/s11219-013-9208-0. URL http://dx.doi.org/10.1007/s11219-013-9208-0.

[7] S. Gulwani. Automating string processing in spreadsheets using input-output examples. In *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '11, pages 317–330, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0490-0. doi: 10.1145/1926385.1926423. URL http://doi.acm.org/10.1145/1926385.1926423.

[8] S. Gulwani, W. R. Harris, and R. Singh. Spreadsheet data manipulation using examples. *Commun. ACM*, 55(8):97–105, Aug. 2012. ISSN 0001-0782. doi: 10.1145/2240236.2240260. URL http://doi.acm.org/10.1145/2240236.2240260.

[9] W. R. Harris and S. Gulwani. Spreadsheet table transformations from examples. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2011, San Jose, CA, USA, June 4-8, 2011*, pages 317–328, 2011. doi: 10.1145/1993498.1993536. URL http://doi.acm.org/10.1145/1993498.1993536.

[10] B. Hartmann, D. MacDougall, J. Brandt, and S. R. Klemmer. What would other programmers do: Suggesting solutions to error messages. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 1019–1028, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-929-9. doi: 10.1145/1753326.1753478. URL http://doi.acm.org/10.1145/1753326.1753478.

[11] T. Lau, S. A. Wolfman, P. Domingos, and D. S. Weld. Programming by demonstration using version space algebra. *Mach. Learn.*, 53(1-2):111–156, Oct. 2003. ISSN 0885-6125. doi: 10.1023/A:1025671410623. URL http://dx.doi.org/10.1023/A:1025671410623.

[12] V. Le and S. Gulwani. Flashextract: A framework for data extraction by examples. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '14, pages 542–553, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2784-8. doi: 10.1145/2594291.2594333. URL http://doi.acm.org/10.1145/2594291.2594333.

[13] C. Le Goues, M. Dewey-Vogt, S. Forrest, and W. Weimer. A systematic study of automated program repair: Fixing 55 out of 105 bugs for $8 each. In *Proceedings of the 34th International Conference on Software Engineering*, ICSE '12, pages 3–13, Piscataway, NJ, USA, 2012. IEEE Press. ISBN 978-1-4673-1067-3. URL http://dl.acm.org/citation.cfm?id=2337223.2337225.

[14] A. Leung, J. Sarracino, and S. Lerner. Interactive parser synthesis by example. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation, Portland, OR, USA, June 15-17, 2015*, pages 565–574, 2015. doi: 10.1145/2737924.2738002. URL http://doi.acm.org/10.1145/2737924.2738002.

[15] H. Lieberman. *Your wish is my command: Programming by example*. Morgan Kaufmann, 2001.

[16] F. Long and M. Rinard. Automatic patch generation by learning correct code. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL 2016, pages 298–312, 2016.

[17] T. M. Mitchell. Generalization as search. *Artif. Intell.*, 18(2), 1982.

[18] P. Osera and S. Zdancewic. Type-and-example-directed program synthesis. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation, Portland, OR, USA, June 15-17, 2015*, pages 619–630, 2015. doi: 10.1145/2737924.2738007. URL http://doi.acm.org/10.1145/2737924.2738007.

[19] G. Piatetsky-Shapiro. Discovery, analysis and presentation of strong rules. In G. Piatetsky-Shapiro and W. J. Frawley, editors, *Knowledge Discovery in Databases*, pages 229–248. AAAI Press, 1991.

[20] V. Raychev, M. Schäfer, M. Sridharan, and M. T. Vechev. Refactoring with synthesis. In *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2013, part of SPLASH 2013, Indianapolis, IN, USA, October 26-31, 2013*, pages 339–354, 2013. doi: 10.1145/2509136.2509544. URL http://doi.acm.org/10.1145/2509136.2509544.

[21] V. Raychev, M. Vechev, and E. Yahav. Code completion with statistical language models. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '14, pages 419–428, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2784-8. doi: 10.1145/2594291.2594321. URL http://doi.acm.org/10.1145/2594291.2594321.

[22] V. Raychev, M. T. Vechev, and A. Krause. Predicting program properties from "big code". In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, pages 111–124, 2015. doi: 10.1145/2676726.2677009. URL http://doi.acm.org/10.1145/2676726.2677009.

[23] E. Schkufza, R. Sharma, and A. Aiken. Stochastic superoptimization. In *Architectural Support for Programming Languages and Operating Systems, ASPLOS '13, Houston, TX, USA - March 16 - 20, 2013*, pages 305–316, 2013. doi: 10.1145/2451116.2451150. URL http://doi.acm.org/10.1145/2451116.2451150.

[24] R. Singh and S. Gulwani. Synthesizing number transformations from input-output examples. In *Proceedings of the 24th International Conference on Computer Aided Verification*, CAV'12, pages 634–651, Berlin, Heidelberg, 2012. Springer-Verlag. ISBN 978-3-642-31423-0. doi: 10.1007/978-3-642-31424-7_44. URL http://dx.doi.org/10.1007/978-3-642-31424-7_44.

[25] R. Singh and S. Gulwani. Learning semantic string transformations from examples. *Proc. VLDB Endow.*, 5(8):740–751, Apr. 2012. ISSN 2150-8097. doi: 10.14778/2212351.2212356. URL http://dx.doi.org/10.14778/2212351.2212356.

[26] R. Singh and A. Solar-Lezama. Synthesizing data structure manipulations from storyboards. In *SIGSOFT FSE*, pages 289–299, 2011.

[27] A. Solar-Lezama. *Program Synthesis By Sketching*. PhD thesis, EECS Dept., UC Berkeley, 2008.

[28] A. Solar-Lezama, R. Rabbah, R. Bodik, and K. Ebcioglu. Programming by sketching for bit-streaming programs. In *PLDI*, 2005.

[29] A. Udupa, A. Raghavan, J. V. Deshmukh, S. Mador-Haim, M. M. K. Martin, and R. Alur. TRANSIT: specifying protocols with concolic

snippets. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '13, Seattle, WA, USA, June 16-19, 2013*, pages 287–296, 2013. doi: 10.1145/2462156.2462174. URL http://doi.acm.org/10.1145/2462156.2462174.

[30] M. Yakout, A. K. Elmagarmid, J. Neville, M. Ouzzani, and I. F. Ilyas. Guided data repair. *Proc. VLDB Endow.*, 4(5):279–289, Feb. 2011. ISSN 2150-8097. doi: 10.14778/1952376.1952378. URL http://dx.doi.org/10.14778/1952376.1952378.

[31] Y. Yuan, R. Alur, and B. T. Loo. Netegg: Programming network policies by examples. In *Proceedings of the 13th ACM Workshop on Hot Topics in Networks, HotNets-XIII, Los Angeles, CA, USA, October 27-28, 2014*, pages 20:1–20:7, 2014. doi: 10.1145/2670518.2673879. URL http://doi.acm.org/10.1145/2670518.2673879.

A. Proofs of Theorems 1, 2, and 3

We first define a notion of completeness for a symbolic rule. Intuitively a symbolic rule has to summarize all possible correct rules to be complete.

Definition 6 (Command-string completeness)**.** *Let*

$$r = \textbf{\textit{match}}\ cmd\ \textbf{\textit{and}}\ err\ \rightarrow_s\ fixes$$

be a symbolic rule such that $cmd = [c_1, \ldots, c_a]$, $err = [m_1, \ldots, m_b]$, $fixes = [f_1, \ldots, f_c]$, *and* $E = [e_1, \ldots, e_n]$ *a sequence of examples. We say that* cmd *is complete for* E *and produces variables* V_1, $CompC(cmd, E, V_1)$ *iff for every* $1 \leq i \leq a$*:*

- $c_i = Str(s)$ *(for some* s*) iff every example* $(\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})$ *in* E *is such that* $\bar{s}_{cmd} = [s_1, \ldots, s_k]$*, for some* k*, and* $s_j = s$*.*
- $c_i = Var\text{-}Match(j, l, r)$ *iff* $i \in V$*,* $i = j$*, and there exists two examples* $(\bar{s}^1_{cmd}, \bar{s}^1_{err}, \bar{s}^1_{fix})$ *and* $(\bar{s}^2_{cmd}, \bar{s}^2_{err}, \bar{s}^2_{fix})$ *in* E *such that* $\bar{s}^1_{cmd} = [s_1, \ldots, s_k]$*,* $\bar{s}^2_{cmd} = [s'_1, \ldots, s'_k]$*, for some* k*, and* $s_j \neq s'_j$*.*
- *If* $c_i = Var\text{-}Match(j, l, r)$*, then for every example* $(\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})$ *in* E *where* $\bar{s}_{cmd} = [s_1, \ldots, s_k]$*,* l *is a prefix of* s_j*. Moreover,* l *is the longest such prefix.*
- *If* $c_i = Var\text{-}Match(j, l, r)$*, then for every example* $(\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})$ *in* E *where* $\bar{s}_{cmd} = [s_1, \ldots, s_k]$*,* r *is a suffix of* s_j*. Moreover,* r *is the longest such suffix.*

Definition 7 (Error-string completeness)**.** *Analogously, we say that* err *is complete for* E *and produces variables* V_2*,* $CompE(err, E, V_2)$*, iff for every* $1 \leq i \leq b$*:*

- $m_i = Str(s)$ *(for some* s*) iff every example* $(\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})$ *in* E *is such that* $\bar{s}_{cmd} = [s_1, \ldots, s_k]$*, for some* k*, and* $s_j = s$*.*
- $m_i = Var\text{-}Match(j, l, r)$ *iff* $i \in V$*,* $i + a = j$*, and there exists two examples* $(\bar{s}^1_{cmd}, \bar{s}^1_{err}, \bar{s}^1_{fix})$ *and* $(\bar{s}^2_{cmd}, \bar{s}^2_{err}, \bar{s}^2_{fix})$ *in* E *such that* $\bar{s}^1_{err} = [s_1, \ldots, s_k]$*,* $\bar{s}^2_{err} = [s'_1, \ldots, s'_k]$*, for some* k*, and* $s_i \neq s'_i$*.*
- *If* $m_i = Var\text{-}Match(j, l, r)$*, and* $i + a = j$*, then for every example* $(\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})$ *in* E *where* $\bar{s}_{err} = [s_1, \ldots, s_k]$*,* l *is a prefix of* s_i*. Moreover,* l *is the longest such prefix.*
- *If* $m_i = Var\text{-}Match(j, l, r)$*, and* $i + a = j$ *then for every example* $(\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})$ *in* E *where* $\bar{s}_{err} = [s_1, \ldots, s_k]$*,* r *is a suffix of* s_i*. Moreover,* r *is the longest such suffix.*

Definition 8 (Input completeness)**.** *If both* $CompC(cmd, E, V_1)$ *and* $CompE(err, E, V_2)$ *hold we say that* cmd *and* err *are complete for* E *and produce variables* $V_1 \cup V_2$*,* $Comp(cmd, err, E, V_1 \cup V_2)$*.*

Definition 9 (Partial f_i-completeness)**.** *We say that* f_i *is partially complete,* $PCompFi(f_i, E, i)$*, with respect to* E *if the following condition holds:* $f_i = [\text{Fstr}(s)]$ *(for some* s*) iff every example* $(\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})$ *in* E *is such that* $\bar{s}_{fix} = [s_1, \ldots, s_k]$*, for some* k*, and* $s_i = s$*. If for every* $1 \leq i \leq c$*,* $PCompFi(f_i, E, i)$ *holds, then we say that that* $fixes$ *is partially complete with respect to* V*,* $PCompF(fixes, E)$*.*

Definition 10 (f_i-completeness)**.** *If there exists a* $V = \{i_1, \ldots, i_j\}$ *such that* $Comp(cmd, err, E, V)$ *we say that* f_i *is complete with respect to* V *at position* i*,* $CompFi(f_i, E, V, i)$*, iff:*

- $PCompFi(f_i, E, i)$*.*
- $f_i = [t_1, \ldots, t_m]$ *such that for all* ind*,* $t_{ind} \neq \text{Fstr}(s)$ *(for any* s*) iff the following properties hold.*
  - *For every* $ind \leq m$ *and* $l \leq k$*,* $[\![t_{ind}]\!]^{\textbf{\textit{fun}}}_{\sigma_l} = \bar{s}^l_{fix}[i]$ *(where for all* $y \leq j$*,* $\sigma_l(i_y) = (\bar{s}^l_{cmd} @ \bar{s}^l_{err})[i_y]$*).*
  - *If there exists a* Sub-lr *function* t *such that for every* $ind \leq n$ *and* $l \leq k$*, ,* $[\![t]\!]^{\textbf{\textit{fun}}}_{\sigma_l} = \bar{s}^l_{fix}[i]$ *(where for all* $y \leq j$*,* $\sigma_l(i_y) = (\bar{s}^l_{cmd} @ \bar{s}^l_{err})[i_y]$*), then there exists* $x \leq m$ *such that* $t_x = t$*.*

Definition 11 (Fix completeness)**.** *If for every* $1 \leq i \leq c$*,* $CompFi(f_i, E, V, i)$ *holds, then we say that that* $fixes$ *is complete with respect to* V*,* $CompF(fixes, E, V)$*.*

Definition 12 (Rule completeness)**.** *If there is a* V *such that* $Comp(cmd, err, E, V)$ *and* $CompF(fixes, E, V)$*, we say that* r *is complete for* E*,* $CompR(r, E)$*.*

Notice that for any permutation E' of E $CompR(r, E)$ iff $CompR(r, E')$.

Proposition 13 ($CompR$ and `con`)**.** *Let* E *a sequence of examples and* $r = \textbf{\textit{match}}\ cmd\ \textbf{\textit{and}}\ err\ \rightarrow_s\ fixes$ *be a symbolic rule. If* $CompR(r, E)$ *then every concrete rule* $r' \in \texttt{con}(r)$ *is consistent with any example in* E*. Moreover, for every (non-symbolic) rule*

$$r_1 = \textbf{\textit{match}}\ cmd\ \textbf{\textit{and}}\ err\ \rightarrow\ [t_1, \ldots, t_n]$$

consistent with $E = [(\bar{s}^1_{cmd}, \bar{s}^1_{err}, \bar{s}^1_{fix}), \ldots, (\bar{s}^n_{cmd}, \bar{s}^n_{err}, \bar{s}^n_{fix})]$ *the following is true: if for every* i*,* $t_i \neq \text{Fstr}(s)$ *iff for some* $j_1, j_2 \leq n$ $\bar{s}^{j_1}_{fix}[i] \neq \bar{s}^{j_2}_{fix}[i]$*, then* $r_1 \in \texttt{con}(r)$*.*

Proof. Immediate from Definition 12 and the definition of `con`.

More specifically, assume there is some $r' \in \texttt{con}(r)$ not consistent with an example $e \in E$. The only way it can be inconsistent with $\bar{s}_{cmd}$ or $\bar{s}_{err}$ is if some match expression in cmd or err is a fixed string not equal to some string in $\bar{s}_{cmd}$ or $\bar{s}_{err}$. Definition 8 precludes this possibility.

Similarly, it follows from Definition 9 that if one of the fix expressions is a constant string, then it is consistent with every example fix. By Definition 10, if it is a set of SUB-LR expressions, then each one is consistent with every example.

□

We now show that SYNTHSUBSTRING and SYNTHFIX have the intended behaviour with respect to Definition 12.

Lemma 14 (Correctness of SYNTHSUBSTRING)**.** *Let* $E = [(\bar{s}^1_{cmd}, \bar{s}^1_{err}, \bar{s}^1_{fix}), \ldots, (\bar{s}^n_{cmd}, \bar{s}^n_{err}, \bar{s}^n_{fix})]$ *be a sequence of examples where every* $\bar{s}^v_{fix}$ *has length* m*,* $V = \{i_1, \ldots, i_j\}$ *be a set of variables, and* $i \leq m$ *be an index in the output fix. If there exist* a *and* b *such that* $\bar{s}^a_{fix}[i] \neq \bar{s}^b_{fix}[i]$*, then* $\text{SynthSubstring}(E, V, i) = S$ *iff* $Comp(S, E, V, i)$*.*

Proof. This lemma states that SYNTHSUBSTRING returns all the SUB-LR functions consistent with the given examples. Again the proof is by induction on the length of E. The base case $|E| = 1$ follows from the definition of ALLSUBSTRINGS. The inductive step is also trivial: SYNTHSUBSTRING simply runs all the functions computed so far on the added example and filters out those that are not consistent with it. Since, by IH, the varaible F was correct at

the beginning of the loop, it remains correct. Notice that the order of the examples does not matter. □

Lemma 15 (Correctness of SYNTHFIX)**.** *Let $S = [s_1, \ldots, s_n]$ be a list of strings, $T = [t_1, \ldots, t_n]$ be a sequence of symbolic fix expressions,*

$$E = [(\bar{s}^1_{cmd}, \bar{s}^1_{err}, \bar{s}^1_{fix}), \ldots, (\bar{s}^n_{cmd}, \bar{s}^n_{err}, \bar{s}^n_{fix})]$$

be a sequence of examples where every $\bar{s}^v_{fix}$ has length m, $e = (\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix})$ be an example such that for each $1 \leq i \leq n$, $\bar{s}_{fix}[i] = s_i$, and V be a set of variables. If $PCompF(T, E)$, then SynthFix$(S, T, e :: E, V) = fixes$ *iff* $CompF(fixes, e :: E, V)$.

Proof. Immediate by induction on m and by using Lemma 14 at each step.
As SYNTHFIX iterates over each string in the example fix, and compares it to the corresponding fix expression, we can see that the only interesting case of the proof is when t_i is either a fixed string not equal to the correpsonding $\bar{s}^{n+1}_{fix}[i]$ in e, or t_i is a set of SUB-LR expressions. In this case, correctness follows directly from Lemma 14, as SYNTHFIX calls SYNTHSUBSTRING to refine the set of substring operation to those generating $\bar{s}^{n+1}_{fix}[i]$ as well as all other $\bar{s}_{fix}[i]$ values.

□

Next we show the correctness of REFINERULE and FINDVARIABLES.

Lemma 16 (Correctness of FINDVARIABLES)**.** *Let $E = [(\bar{s}^1_{cmd}, \bar{s}^1_{err}, \bar{s}^1_{fix}), \ldots, (\bar{s}^n_{cmd}, \bar{s}^n_{err}, \bar{s}^n_{fix})]$ be a sequence of examples, and cmd and err be two sequences of match expressions of lengths $|\bar{s}^i_{cmd}|$ and $|\bar{s}^i_{err}|$ respectively. If there exists two sets of variables V_1 and V_2 such that $CompC(cmd, E[1 :: (n-1)], V_1)$ and $CompE(err, E[1 :: (n-1)], V_2)$, then*

- *if* $(cmd', V_1') =$ FindVariables$(\bar{s}^n_{cmd}, cmd, 0)$, *then* $CompC(cmd', E, V_1')$.
- *if* $(err', V_2') =$ FindVariables$(\bar{s}^n_{err}, err, |\bar{s}^i_{cmd}|)$, *then* $CompE(err', E, V_2')$.

Proof. The two statements can be proved separately but the proof is identical. The proofs are both by induction on the length of the first argument of FINDVARIABLES and are very simple case analysis following Definitions 6 and 7.

W.L.O.G, consider the case for cmd. Over a run of FINDVARIABLES, a STR(s) in cmd is only changed if, at the nth iteration of the loop, $cmd[i]$ is STR(S) where S is not equal to $\bar{s}^n_{cmd}[i]$. In that case FINDVARIABLES introduces a new variable match expression VAR-MATCH(i, l, r) where l and r are (respectively) the longest shared prefix and suffix of S and $\bar{s}^n_{cmd}[i]$. By the I.H, cmd was command complete. It follows directly from Definition 6 that command completeness continues to hold.

A VAR-MATCH(i, l, r) in cmd is only changed if, at iteration n, the longest common prefix (suffix), x, shared by $\bar{s}^n_{cmd}[i]$ and l (r) is not equal to l (r). In this case VAR-MATCH(i, l, r) is replaced by VAR-MATCH(i, l', r) (VAR-MATCH(i, l, r')), where r' (l') is equal to the common prefix (suffix) x. By the I.H, l was the longest prefix (suffix) shared by the first $n-1$ values of $\bar{s}^k_{cmd}[i]$. Clearly, $x = r'$ (l') is the longest prefix (suffix) shared by the first n values of $\bar{s}^k_{cmd}[i]$. Moreover, by the I.H. cmd satisfied all other criteria for command completeness. Thus, it follows that command completeness continues to hold. It is clear that determining the shorteset prefix and suffix shared by every $\bar{s}^k_{cmd}[i]$ does not depend on the order in which the examples are presented.

□

Lemma 17 (Correctness of REFINERULE)**.** *Let r be a symbolic rule, E be a sequence of examples, e be an example, and $r' =$* RefineRule(r, E, e). *If $CompR(r, E)$ then $CompR(r', e :: E)$.*

Proof. Immediate from Definition 12 and Lemmas 15 and 16. □

Lemma 18 (Correctness of SYNTHRULES)**.** *Let E be a sequence of examples. If $r' =$* SynthRules(E) *then $CompR(r, E)$.*

Proof. By induction on the length of E. Case $|E| = 1$ and $E = [e]$: the result of CONSTRULE(e) clearly satisfies $CompR(r, E)$. The inductive step follows from Lemma 17. □

We can now conclude the proofs of Theorems 1, 2, and 3. Theorem 1 and 2 follow from the order invariance of $CompR(r, E)$ and from Proposition 13 and Lemmas 18.

Proof of Theorem 3 Let

$$E = [(\bar{s}^1_{cmd}, \bar{s}^1_{err}, \bar{s}^1_{fix}), \ldots, (\bar{s}^n_{cmd}, \bar{s}^n_{err}, \bar{s}^n_{fix})]$$

be a sequence of examples,

$$\text{SYNTHRULES}(E) = r = \textbf{match } cmd \textbf{ and } err \rightarrow_s fixes$$

be the symbolic rule for E and

$$r' = \textbf{match } cmd' \textbf{ and } err' \rightarrow fix'$$

be a concrete rule that is consistent with all the examples in E. If r' belongs to $\texttt{con}(r)$ then we are done. Assume it doesn't.

Case 1. $cmd = cmd'$ and $err = err'$. Then by Proposition 13 there exists some $t_{i_1}, \ldots, t_{i_k}$ such that each t_{i_v} is of the form SUB-LR$(p^v_L, p^v_R, l^v, r^v,$ VAR$(j_v))$, and for all $j \leq n$ $\bar{s}^j_{fix}[i_v] = s$ for some s (i.e., the output is a function of the input, but all examples can be captured using a constant output). In this case it is enough to create a new example e' starting from any example $(\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix}) \in E$ where for each i_v we modify $s_{j_v} = (\bar{s}_{cmd} @ \bar{s}_{err})[j_v]$ so that SUB-LR$(p^v_L, p^v_R, l^v, r^v,$ VAR$(j_v))$ now returns a value different from the previous one (this can be done by simply adding a new character between p^v_L and p^v_R). $\bar{s}_{fix}$ is replaced by the result of applying r' to the modified input.

Case 2. $cmd@err \neq cmd'@err'$. This means that r' uses more variables then r (notice that from Definition 12 the set of variables used by r' is necessary). This can be fixed by changing the input of any example $(\bar{s}_{cmd}, \bar{s}_{err}, \bar{s}_{fix}) \in E$. For each variable VAR$(i)$ that is in $cmd'@err'$ but not in $cmd@err$ replace the string $(\bar{s}_{cmd} @ \bar{s}_{err})[i] = a_1 \ldots a_n$ with $ba_1 \ldots a_n$ where b is a symbol not appearing in $a_1 \ldots a_n$. $\bar{s}_{fix}$ is replaced by the result of applying r' to the modified input. If the rule is still not in $\texttt{con}(r')$ we can then modify the example using the techniques from Case 1 since now $cmd = cmd'$ and $err = err'$.